The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XLVIII-2/W9-2025
ISPRS Intl. Workshop “Photogrammetric and computer vision techniques for environmental and infraStructure monitoring, Biometrics and Biomedicine”
PSBB25 , 9–11 June 2025, Moscow, Russia

Using Uncertainty to Expand Training Sets for Mineral Segmentation in Geological Images

Olesya Indychko[1], Dmitry Korshunov[2], Alexander Khvostikov[1]

[1] Faculty of Computational Mathematics and Cybernetics, Lomonosov Moscow State University, Moscow, Russia
olesyaindychko@gmail.com, khvostikov@cs.msu.ru
[2] Geological Institute, Russian Academy of Sciences, Moscow, Russia
d.korshunov@ginras.ru

Keywords: Color Adaptation of Images, Polished Sections, Color Correction, Mineral Segmentation, Deep Learning

Abstract

This paper introduces a method to enhance the process of creating training datasets for mineral segmentation in geological images by addressing the challenges posed by color distortion. Such distortions, which stem from differing imaging equipment, create inconsistencies in color and brightness that hinder effective segmentation by neural networks. By utilizing the hyperbolic active learning method (HALO), the proposed method targets regions with epistemic uncertainty in neural network models, enabling the focused expansion of training data instead of revisiting entire images. Experiments were conducted on the LumenStone S1v2 dataset, revealing that the hyperbolic radius effectively correlates with error maps, thereby highlighting uncertain regions that need annotation. This method significantly reduces the manual effort needed from specialists during annotation and promises to improve segmentation accuracy. Future developments include integrating these techniques into a complete neural network pipeline, leveraging color correction and uncertainty mapping for more precise mineral segmentation.

1. Introduction

Mineral segmentation in geological images is a fundamental task in modern geoscience. This process is important for different applications including ore deposit characterization, petrological studies, and mineral resource assessment. However, achieving reliable segmentation results remains challenging due to variations in sample preparation and inconsistencies in imaging conditions, which has motivated the development of specialized computational approaches to maintain accuracy across different datasets.

We propose a method for automatically identifying missing mineral information when segmenting images of geological sections obtained using different imaging equipment. The goal of this work is to enhance the performance of a segmentation neural network model by strategically expanding the training dataset.

A key challenge in automatic mineral segmentation is the nonhomogeneity between training and test data, which arises due to variations in imaging conditions, equipment, and sample preparation. This inconsistency can degrade model performance when applied to new, unseen data. The example of nonhomogenity between geological images is presented in Fig. 1. This images were specifically taken using different equipment and light conditions.

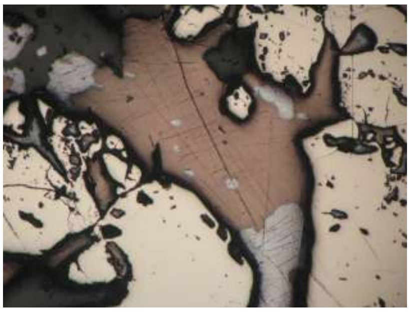

a)

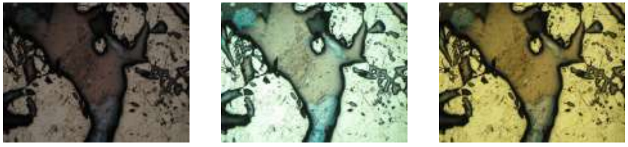

b) c) d)

Figure 1. Example of nonhomogenity of geological images.
a) — Image from the train subset of LumenStone S1v2 dataset,
b), c), d) — Images of the same mineral area captured under different conditions

To address this issue, two primary approaches have been explored.

The approach of **model adaptation to diverse input data** focuses on improving the robustness of the neural network itself, enabling it to generalize across heterogeneous datasets. Techniques include domain adaptation (Wang and Deng, 2018), where the model is fine-tuned to align feature distributions between source and target domains, and the use of data augmentation to simulate variations in lighting, resolution, and color balance. Recent advances in self-supervised learning (Chen et al., 2020b) and contrastive learning (Chen et al., 2020a) have also shown promise in learning invariant representations from unlabeled geological images. However, these methods often require extensive computational resources and may still struggle with extreme domain shifts.

Another common strategy to improve model robustness is data augmentation, where synthetic variations (e.g., rotations, flips, and noise) are introduced to simulate real-world heterogeneity. However, in mineral segmentation, standard augmentations often fail to capture the complex color and brightness distortions caused by different microscopes, lighting conditions, or sample preparations. While advanced techniques like generative adversarial networks (GANs) (Goodfellow et al., 2020) or neural style transfer (Jing et al., 2020) can simulate domain shifts, they require large amounts of paired data and may produce unrealistic artifacts. As a result, augmentation alone is insufficient to bridge the gap between highly dissimilar datasets.

This limitation motivates an alternative solution: **expanding the training set with partially annotated real images** that reflect true domain variations. Instead of relying on synthetic distortions, our method identifies and incorporates uncertain regions (areas where the model's predictions are least confident) into the training process. This ensures that new samples address the model's weaknesses while maintaining realistic data distributions.

Instead of modifying the model, the strategy of **input image adaptation to a reference dataset** transforms input images to match the characteristics of a reference dataset used during training. An example of this approach is the color correction algorithm (Indychko et al., 2023), which aligns the color and brightness of new images with those of polished sections used for model training. The method leverages a transition matrix between color spaces (Cohen, 1988) to minimize inconsistency caused by different imaging conditions. While being effective, this approach assumes that color distribution is the primary source of variation and may not account for other forms of heterogeneity, such as texture or scale differences.

To further improve mineral segmentation accuracy, we suggest **combining color correction with an intelligent training set expansion strategy**. Since manual annotation of mineral samples is time-consuming and requires expert knowledge, it is crucial to prioritize only the most informative regions for labeling. Inspired by recent advances in active learning, we propose an automated approach to identify uncertain regions where the model's predictions are least confident. Specifically, we adapt the Hyperbolic Active Learning (HALO) framework (Franco et al., 2024) to geological image segmentation, enabling efficient selection of samples that maximize model improvement with minimal human intervention.

By integrating uncertainty-aware sample selection with domain adaptation techniques, our method aims to enhance mineral segmentation accuracy while reducing annotation effort. This work contributes to the broader goal of automating geological analysis, making it more scalable and accessible for resource exploration and material science applications.

2. Used Data and Models

In this work we used the S1v2 subset of the LumenStone dataset[1] containing 3396×2547 pixel photos of polished sections with full semantic segmentation masks for each image. An example of image and its semantic segmentation mask from train subset of LumenStone S1v2 dataset are presented in Fig. 2.

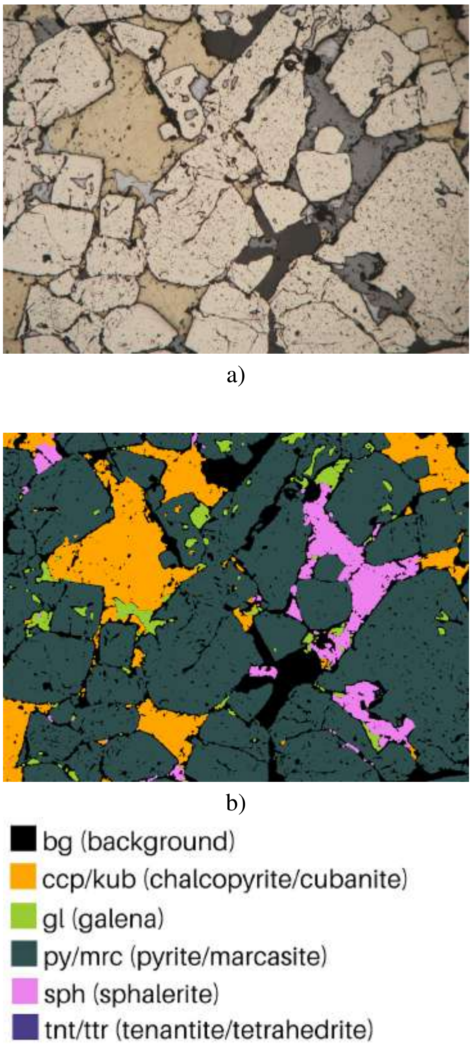

a)

b)

c)

Figure 2. Example of data from S1v2 train subset of LumenStone dataset.
a) — Image from the train subset of LumenStone S1v2 dataset,
b) — Semantic segmentation mask for the image a),
c) — Color legend for segmentation mask b).

During the actual work of geologists, using a trained segmentation model, it was noticed that color-brightness distortions occur in images, depending on the equipment used for imaging. We simulate the distortions by using the following equipment:

- Carl Zeiss AxioScop 40 with a Canon Powershot g10 camera,
- Carl Zeiss AxioImager m1 with a Canon EOS D400 camera,
- LOMO plm 215 with a Canon EOS D400 camera.

To systematically study these real-world color and lightning distortions, we created experimental datasets that mimic the variations observed in practice. Based on the S1v2 test set, three datasets with different color-brightness distortions were created, corresponding to the real data. Next, we will refer to these datasets as the "blue", "pink", and "yellow" distortions. The example of simulated image distortions is shown in Fig. 3.

---

[1] `https://imaging.cs.msu.ru/en/research/geology/lumenstone`

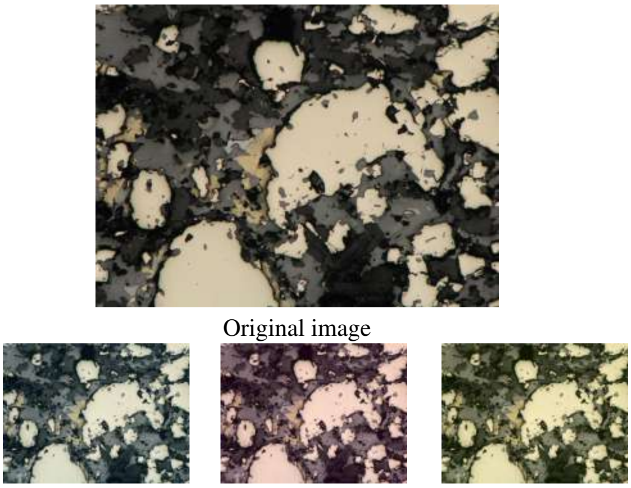
Original image
”Blue” distortion ”Pink” distortion ”Yellow” distortion

Figure 3. Three types of distortion for image from LumenStone S1v2 test subset.

Each distortion dataset comprises 10 gradually intensifying variations for every image in the original test set. These variations were generated by incrementally applying color channel adjustments over 10 sequential steps, thus producing a systematic progression of distortion for each source image.

For the ”yellow” distortion, the RGB channels were modified with intensity values of (R:+15, G:+20, B:+2), creating a yellowish tint by predominantly enhancing the red and green components.

The ”blue” distortion was implemented with RGB adjustments of (R:+2, G:+15, B:+20), resulting in a bluish appearance through stronger enhancement of the green and blue channels.

Similarly, the ”pink” distortion utilized RGB modifications of (R:+20, G:+2, B:+15), producing a pinkish hue by amplifying primarily the red and blue components.

These color modifications were applied incrementally across 10 steps, with each subsequent image receiving an additional fraction of the total intensity adjustment. This methodical approach allowed us to simulate a spectrum of realistic color variations that might be encountered in practical geological imaging scenarios while maintaining control over the distortion parameters.

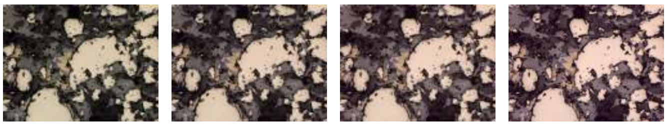
Original image 4th step 7th step 10th step

Figure 4. Example of increasing ”pink” color distortion degree.

In this work we used two models as trained segmentation models: ResUnet (Khvostikov et al., 2021) and customized PSPNet with dilated convolutions and ResNet18 encoder. Both models were trained on the original 64 images of LumenStone S1v2 train subset.

3. Proposed Method

Recently it was proposed to narrow the search for regions for a partial segmentation mask to areas of epistemic uncertainty (Depeweg et al., 2018) in a neural network model. To search for areas of uncertainty, in this paper we adapt the hyperbolic active learning method (HALO) (Franco et al., 2024).

A method for calculating the hyperbolic radius is proposed in the paper (GhadimiAtigh et al., 2022). To calculate the hyperbolic radius, we need to operate in the Poincaré ball hyperbolic manifold, defined as the pair $(\mathbb{D}^N, g^{\mathbb{D}})$, where $\mathbb{D}^N = \{x \in \mathbb{R}^N : \|x\| < 1\}$ is the manifold and $g^{\mathbb{D}}_x = (\lambda_x)^2 g^{\mathbb{E}}$ is the associated Riemannian metric, $\lambda_x = \dfrac{2}{1 - \|x\|^2}$ is the conformal factor and $g^{\mathbb{E}} = \mathbb{I}^N$ is the Euclidean metric tensor. Neural networks first extract a feature vector v in Euclidean space, which is subsequently projected into the Poincaré ball via exponential map:

$$\exp_x(v) = x \oplus \left(\frac{v}{\|v\|} \tanh \left(\frac{\lambda_x \|v\|}{2} \right) \right), \tag{1}$$

where $x \in \mathbb{D}^N$ is the anchor and $\oplus$ is the Möbius hyperbolic addition. The latter is defined as

$$h \oplus w = \frac{(1 + 2\langle h, w\rangle + \|w\|^2)v + (1 - \|h\|^2)w}{1 + 2\langle h, w\rangle + \|h\|^2\|w\|^2}, \tag{2}$$

where h and w are hyperbolic vectors.

The hyperbolic radius is the Poincaré distance from the origin of a ball to the projection $h \in \mathbb{D}^N$ of a feature vector v. If $h \in \mathbb{D}^N$ is the projection of v, the hyperbolic radius is defined as

$$d(h, 0) = 2 \tanh^{-1}(\|h\|). \tag{3}$$

Pixel-by-pixel calculation of the hyperbolic radius allows us to create a hyperbolic radius heat map. The example of hyperbolic radius visualisation as a heatmap for the image from Fig. 5 is shown in Fig. 6.

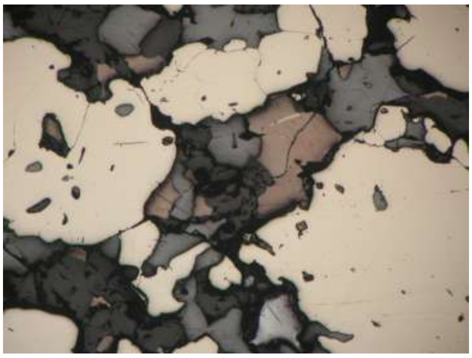

Figure 5. Image from the test set of LumenStone S1v2 test subset.

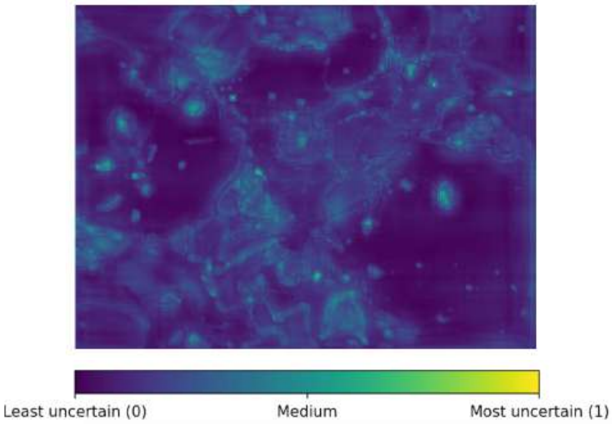

Figure 6. Hyperbolic radius map for the image from Fig. 5.

We hypothesize that large values of the radius correspond to low confidence in the model. To test this, we created heatmaps of segmentation model errors. These error maps are calculated by taking the degree of distortion for each pixel in the image and assigning it a value based on how fast the trained model is making mistakes when distorting the image. The faster the model makes mistakes, the larger the value for that pixel in the error map will be.

Let M_i be the binary error mask at distortion degree i, n — the total number of distortions (in our expirements $n = 10$), H — the resulting heatmap. Therefore, the error heatmap can be expressed as:

$$H = \sum_{i=0}^{n-1} (n - i) \cdot (1 - M_i). \tag{4}$$

The example of error heatmap for the image from Fig. 5 is shown in Fig. 7. The color scheme represents on what step of consequent image distortion trained segmentation model makes mistake (fails).

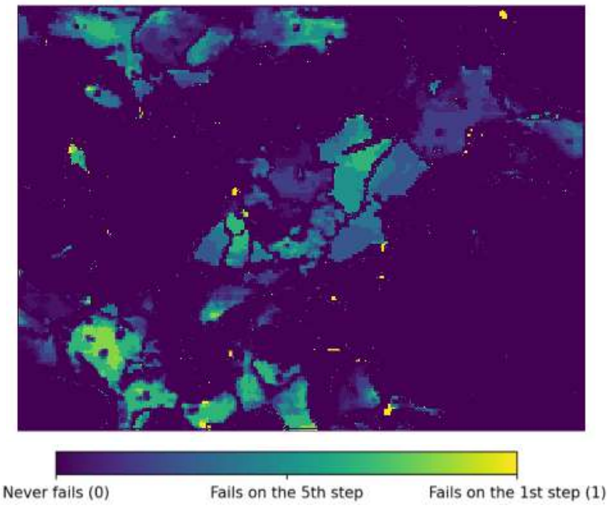

Figure 7. Error map for the image from Fig. 5.

To confirm our hypothesis, we calculated the weighted correlation $r_w(X, Y)$ between the hyperbolic radius (X) and the error maps (Y) with the weigths $w = Y$ for each mineral from the dataset separately. We specifically chose weighted correlation rather than standard correlation to penalize our uncertainty detection method only in cases where the segmentation model made errors but the uncertainty maps failed to highlight those pixel areas. By using the error maps as weights ($w = Y$), we give greater importance to pixels where errors actually occur, ensuring that our evaluation focuses on the method's ability to identify problematic regions rather than its performance on easily classified pixels.

$$r_w = \frac{\sum w_i (X_i - \bar{X}_w)(Y_i - \bar{Y}_w)}{\sqrt{\left(\sum w_i (X_i - \bar{X}_w)^2\right)\left(\sum w_i (Y_i - \bar{Y}_w)^2\right)}} \tag{5}$$

where $\bar{X}_w = \frac{\sum w_i X_i}{\sum w_i}$ — weighted average of X
X — flattened hyperbolic radius maps,
Y — flattened error maps,
$w = Y$ — weights.

4. Implementation and Results

The proposed method was implemented using Python 3 and pytorch.

The experiments were conducted on two segmentation models, PSPNet and ResUnet, at different stages of training: fully trained (50 epochs), after 10 and 5 epochs respectively. The results suggest that, for the more complex PSPNet model, the hyperbolic radius is more closely correlated with the error maps than for the simpler ResUnet.The results of the experiments on Lumenstone S1v1 subset are presented in Tables 1 and 2.

	Distortion type	ResUnet IOU per mineral				
		Ccp	Gl	Py	Tnt	Brt
Fully-trained	Blue	**0.43**	-0.63	-0.37	**0.20**	**0.22**
	Pink	-0.02	0.00	-0.02	**0.14**	-0.11
	Yellow	-0.48	-0.18	-0.43	0.04	0.07
10 epochs	Blue	**0.33**	-0.63	-0.34	**0.17**	**0.30**
	Pink	-0.10	-0.02	0.00	**0.17**	-0.07
	Yellow	-0.34	-0.16	-0.42	0.05	0.05
5 epochs	Blue	**0.56**	-0.05	-0.09	0.08	-0.20
	Pink	**0.50**	-0.01	0.12	-0.11	**0.19**
	Yellow	0.10	0.04	-0.00	-0.07	0.08

Table 1. Weighted correlation between the error maps and the **hyperbolic radii** for ResUnet model on LumenStone S1v1 Test.

	Distortion type	PSPNet IOU per mineral				
		Ccp	Gl	Py	Tnt	Brt
Fully-trained	Blue	**0.24**	-0.24	**0.58**	**0.41**	**0.62**
	Pink	-0.43	**0.38**	**0.51**	**0.41**	**0.16**
	Yellow	**0.39**	-0.27	-0.17	**0.57**	-0.03
10 epochs	Blue	**0.53**	-0.41	**0.48**	-0.02	**0.62**
	Pink	-0.34	**0.41**	**0.27**	**0.49**	**0.27**
	Yellow	-0.02	-0.32	-0.09	**0.45**	-0.11
5 epochs	Blue	**0.50**	-0.30	-0.03	0.07	**0.49**
	Pink	-0.20	-0.36	**0.32**	**0.49**	0.07
	Yellow	**0.31**	-0.32	-0.17	-0.04	-0.08

Table 2. Weighted correlation between the error maps and the **hyperbolic radii** for PSPNet model on LumenStone S1v1 Test.

Hyperbolic radius maps have been post-processed to identify areas of uncertainty that require further attention from geologists. Threshold filtering, dilation, and erosion were applied to reduce noise and improve the clarity of the map.

As a result, a we receive an image and a map with missing markings. In order to expand the dataset, the specialist only needs to annotate the selected minerals instead of annotating the entire image as before. An example of such data can be seen in Fig. 8. The image shows that the neural network model is uncertain about only one mineral located in the upper right corner. This significantly reduces a specialist's workload.

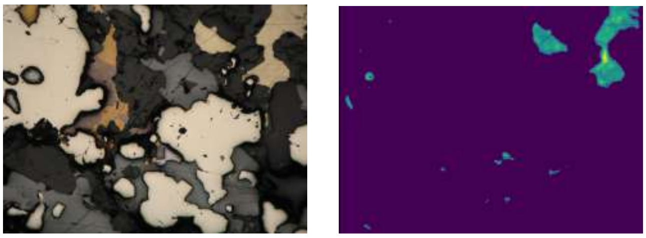
Image
Required annotation

Figure 8. An example of a map with required image annotation from the LumenStone S1v2 dataset.

4.1 Alternative Approaches Comparison

We have considered four different methods for identifying areas of uncertainty. The example of uncertainty maps, obtained by these methods, are presented in Fig. 9. For each method, experiments described in section 3 were conducted. The results for ResUnet and PSPNet models are presented below in the Table 3.

1. **Prediction entropy** — one of the most commonly used approach to uncertainty searching in different machine learning tasks. Such as hydrological data analysis (Chapman, 1986), zero-shot learning in image classification task (Chen et al., 2021), trajectory prediction in autonomous riving (Distelzweig et al., 2025) and so on. Prediction entropy is determined by the following formula

$$H[x] := -\sum_i \mathbb{E}P_i \cdot \log \mathbb{E}P_i, \tag{6}$$

   where $H[x]$ is the Shannon entropy (Shannon, 1948) of the predictive output over the probability $P = (P_1, \dots P_C)$. In our task P_i — is the probability of belonging pixel to i-th segmentation class and C — total number of classes.

2. **Margin sampling** — an approach that quantifies uncertainty by measuring the difference between the probabilities of the two most likely classes (Shin et al., 2021). This metric highlights regions where the model struggles to decisively choose between competing class predictions. Margin sampling can be defined as the difference between the first and the second most probable labels' probabilities

$$P_{marg}[x] := \max \mathbb{E}P_i - \max_2 \mathbb{E}P_i, \tag{7}$$

   where $\max_2$ means the second-largest component. The lower the value of $P_{marg}[x]$, the higher the uncertainty in the model's predictions of pixel x.

   This approach is also widely used in different tasks, including active learning for semantic segmentation (Didari et al., 2024).

3. **Hyperbolic radius** — approach, which was discussed in details in section 3:

$$R_h[x] := d(x_h, 0) = 2\tanh^{-1}(\|h\|), \tag{8}$$

   where x_h is the projection of embedding vector x to a hyperbolic manifold.

   The hyperbolic radius has been studied in detail from various perspectives, including hierarchical representations (Nickel and Kiela, 2017) and hyperbolic vision transformers (Ermolov et al., 2022).

   The results of experiments for hyperbolic radius are presented in Tables 1 and 2.

4. **Production of hyperbolic radius and entropy** — a method, which was originally introduced in HALO approach (Franco et al., 2024). It's pixel-by-pixel multiplication of hyperbolic radius and entropy values:

$$PR_h[x] := R_h[x] \cdot H[x]. \tag{9}$$

Our experiments revealed several important findings about uncertainty estimation for mineral segmentation in geological images:

1. **Model-Specific Performance**. All uncertainty estimation methods showed better results on the larger PSPNet model compared to the smaller ResUnet model. This suggests that model complexity affects the reliability of uncertainty estimation.
2. **Hyperbolic Radius Effectiveness**. The hyperbolic radius consistently performed better than other methods across both models. Its mathematical foundation in hyperbolic space (Franco et al., 2024) allows it to accurately identify uncertain regions at mineral boundaries, even when color and texture vary significantly.
3. **Method Selection Considerations**. While the hyperbolic radius was the best overall method for our datasets, alternative methods also showed good performance with the PSPNet model. This indicates that different uncertainty measures may be suitable for different imaging conditions. The hyperbolic radius can serve as a reliable default method, with other approaches available as alternatives for specific applications.

5. Conclusion

In this paper, we presented a method for automatic identification of areas of uncertainty for a trained segmentation model in order to further enrich the training dataset with images of needed minerals. The method was tested on LumenStone dataset with three varying levels of color distortion, using two different trained segmentation models — ResUnet and PSPNet.

Our comparative analysis of four uncertainty estimation approaches revealed notable differences in performance across model architectures. While all approaches demonstrated satisfactory results when applied to the larger PSPNet architecture, the hyperbolic radius method exhibited slightly better performance. More significantly, when evaluated on the smaller ResUnet architecture, the hyperbolic radius approach substantially

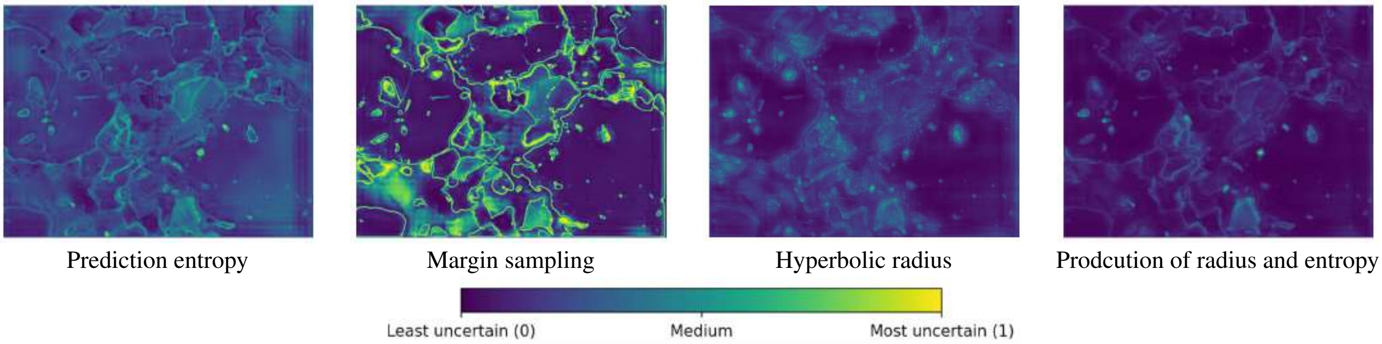
Prediction entropy Margin sampling Hyperbolic radius Prodcution of radius and entropy

Figure 9. Uncertainty maps for image from Fig. 5, obtained by different methods.

Uncertainty identifying	Train stage	Distortion type	ResUnet IOU per mineral					PSPNet IOU per mineral				
			Ccp	Gl	Py	Tnt	Brt	Ccp	Gl	Py	Tnt	Brt
Prediction entropy	Fully-trained	Blue	-0.44	-0.66	-0.09	-0.18	-0.50	-0.06	-0.07	**0.18**	-0.04	**0.52**
		Pink	0.02	0.03	-0.10	-0.26	**0.24**	**0.17**	**0.35**	**0.21**	**0.26**	0.03
		Yellow	**0.46**	0.09	-0.04	-0.02	-0.58	-0.06	-0.05	-0.15	**0.34**	**0.43**
	10 epochs	Blue	-0.39	-0.65	-0.01	-0.26	-0.58	-0.05	-0.15	**0.29**	0.06	**0.50**
		Pink	0.12	-0.02	-0.03	-0.30	**0.22**	**0.34**	**0.18**	0.10	**0.43**	**0.27**
		Yellow	**0.52**	-0.02	-0.17	-0.10	-0.67	0.03	-0.13	-0.11	**0.33**	**0.41**
	5 epochs	Blue	-0.52	-0.34	**0.35**	0.01	-0.61	0.11	-0.05	**0.33**	0.02	**0.20**
		Pink	-0.34	-0.10	-0.25	-0.11	-0.08	**0.17**	0.08	0.14	0.08	-0.01
		Yellow	0.13	**0.25**	**0.40**	0.07	-0.24	0.04	0.14	-0.06	0.10	**0.22**
Margin sampling	Fully-trained	Blue	**0.61**	**0.53**	0.12	**0.58**	**0.54**	**0.24**	**0.72**	**0.39**	**0.57**	**0.59**
		Pink	0.10	0.09	0.13	**0.73**	- 0.34	**0.59**	**0.45**	**0.44**	**0.54**	-0.08
		Yellow	-0.14	-0.27	-0.53	**0.17**	-0.43	-0.20	**0.74**	-0.09	**0.51**	**0.38**
	10 epochs	Blue	**0.62**	**0.48**	0.07	**0.60**	**0.60**	0.14	**0.62**	**0.35**	0.06	**0.56**
		Pink	0.04	**0.19**	0.12	**0.71**	-0.28	**0.31**	**0.34**	**0.27**	**0.58**	-0.17
		Yellow	0.08	**0.24**	**0.34**	**0.67**	**0.63**	-0.04	**0.63**	-0.04	**0.39**	**0.33**
	5 epochs	Blue	**0.61**	**0.23**	**0.41**	**0.50**	**0.64**	0.00	**0.51**	**0.20**	**0.35**	**0.72**
		Pink	**0.38**	0.12	-0.29	**0.60**	-0.09	-0.02	0.00	-0.02	**0.14**	-0.11
		Yellow	-0.33	-0.23	0.04	**0.58**	**0.28**	**0.41**	**0.60**	**0.49**	**0.64**	-0.20
Production of hyperbolic radius and entropy	Fully-trained	Blue	0.12	-0.70	-0.63	0.08	-0.31	**0.23**	-0.37	**0.47**	0.16	**0.61**
		Pink	-0.12	-0.07	-0.03	-0.07	0.04	-0.33	**0.35**	**0.36**	**0.40**	0.15
		Yellow	-0.14	-0.27	-0.53	**0.17**	-0.43	**0.35**	-0.30	-0.18	**0.49**	0.09
	10 epochs	Blue	-0.43	-0.71	-0.68	0.08	-0.15	**0.51**	-0.48	**0.41**	-0.00	**0.61**
		Pink	-0.08	-0.13	0.07	0.05	-0.02	-0.21	**0.32**	**0.18**	**0.44**	**0.26**
		Yellow	0.10	-0.28	-0.58	**0.30**	-0.44	-0.00	-0.35	-0.09	**0.39**	0.00
	5 epochs	Blue	0.02	-0.32	-0.40	**0.20**	-0.61	**0.48**	-0.33	-0.00	0.09	**0.48**
		Pink	0.09	-0.03	**0.27**	**0.17**	0.07	-0.14	-0.34	**0.30**	**0.41**	0.13
		Yellow	0.09	**0.21**	0.04	**0.23**	-0.05	**0.26**	-0.25	-0.11	0.00	-0.06

Table 3. Weighted correlation between the error maps and the uncertainty map, obtained by various methods, for ResUnet and PSPNet models on LumenStone S1v1 Test.

outperformed other methods, highlighting its robustness across model scales.

Our work demonstrates that strategic sample selection guided by uncertainty metrics can effectively identify the most informative regions in geological images, potentially reducing annotation effort while maximizing model improvement. This is particularly valuable in mineral segmentation where expert annotation is time-consuming and costly, and where certain mineral classes or textural features may be underrepresented in training data.

In the future, we plan to expand the training dataset based on the information we receive, as well as to implement a full cycle of neural network operation on images with preliminary color correction, based on partial markings by a geologist on automatically selected areas of uncertainty and integrate it into the `petroscope`[2] python package.

6. Acknowledgments

This work was supported by the Russian Science Foundation, project no. 24-21-00061.

References

Chapman, T. G., 1986. Entropy as a measure of hydrologic data uncertainty and model performance. *Journal of Hydrology*, 85(1), 111-126.

Chen, T., Kornblith, S., Norouzi, M., Hinton, G., 2020a. A Simple Framework for Contrastive Learning of Visual Representations. 119, 1597–1607.

Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G., 2020b. Big self-supervised models are strong semi-supervised learners.

Chen, Z., Huang, Z., Li, J., Zhang, Z., 2021. Entropy-Based Uncertainty Calibration for Generalized Zero-Shot Learning. 139–151.

---

[2] `https://github.com/xubiker/petroscope`

Cohen, J. B., 1988. Color and color mixture: Scalar and vector fundamentals. *Color Research & Application*, 13(1), 5–39.

Depeweg, S., Hernández-Lobato, J. M., Doshi-Velez, F., Udluft, S., 2018. Decomposition of Uncertainty in Bayesian Deep Learning for Efficient and Risk-sensitive Learning. *International Conference on Machine Learning*.

Didari, S., Hu, W., Woo, J. O., Hao, H., Moon, H., Min, S., 2024. Bayesian Active Learning for Semantic Segmentation. *ArXiv*. https://api.semanticscholar.org/CorpusID:271709402.

Distelzweig, A., Look, A., Kosman, E., Janjoš, F., Wagner, J., Valada, A., 2025. Stochasticity in Motion: An Information-Theoretic Approach to Trajectory Prediction. https://arxiv.org/abs/2410.01628.

Ermolov, A., Mirvakhabova, L., Khrulkov, V., Sebe, N., Oseledets, I., 2022. Hyperbolic vision transformers: Combining improvements in metric learning.

Franco, L., Mandica, P., Kallidromitis, K., Guillory, D., Li, Y.-T., Darrell, T., Galasso, F., 2024. Hyperbolic Active Learning for Semantic Segmentation under Domain Shift. *Proceedings of the 41 st International Conference on Machine Learning*.

GhadimiAtigh, M., Schoep, J., Acar, E., van Noord, N., Mettes, P., 2022. Hyperbolic Image Segmentation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4443-4452.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144.

Indychko, O., Khvostikov, A., Korshunov, D., Boguslavskiy, M., Krylov, A., 2023. Color Adaptation in Images of Polished Sections of Geological Specimens. *Computational Mathematics and Modeling*, 33, 487–500.

Jing, Y., Yang, Y., Feng, Z., Ye, J., Yu, Y., Song, M., 2020. Neural Style Transfer: A Review . *IEEE Transactions on Visualization & Computer Graphics*, 26(11), 3365-3385.

Khvostikov, A., Korshunov, D., Krylov, A., Boguslavskiy, M., 2021. Automatic identification of minerals in images of polished sections. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLIV-2/W1-2021, 113-118.

Nickel, M., Kiela, D., 2017. Poincaré Embeddings for Learning Hierarchical Representations. 30.

Shannon, C. E., 1948. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379-423.

Shin, G., Xie, W., Albanie, S., 2021. All you need are a few pixels: semantic segmentation with PixelPick. *IEEE International Conference on Computer Vision Workshops (ICCVW)*, abs/2104.06394.

Wang, M., Deng, W., 2018. Deep visual domain adaptation: A survey. *Neurocomputing*, 312, 135–153.